Acoustic features of objects matched by an echolocating bottlenose dolphin

Caroline M. DeLong[a)] and Whitlow W. L. Au
Hawaii Institute of Marine Biology, P. O. Box 1106, Kailua, Hawaii 96734

David W. Lemonds
Lockheed Martin Orincon, 970 North Kalaheo Avenue, Suite C-215, Kailua, Hawaii 96734

Heidi E. Harley
The Living Seas, Epcot®, Walt Disney World® Resort, New College of Florida, 5700 North Tamiami Trail, Sarasota, Florida 34243

Herbert L. Roitblat
DolphinSearch, Inc., 5855 Olivas Park Drive, Ventura, California 93003

(Received 26 April 2005; revised 27 October 2005; accepted 3 December 2005)

The focus of this study was to investigate how dolphins use acoustic features in returning echolocation signals to discriminate among objects. An echolocating dolphin performed a match-to-sample task with objects that varied in size, shape, material, and texture. After the task was completed, the features of the object echoes were measured (e.g., target strength, peak frequency). The dolphin's error patterns were examined in conjunction with the between-object variation in acoustic features to identify the acoustic features that the dolphin used to discriminate among the objects. The present study explored two hypotheses regarding the way dolphins use acoustic information in echoes: (1) use of a single feature, or (2) use of a linear combination of multiple features. The results suggested that dolphins do not use a single feature across all object sets or a linear combination of six echo features. Five features appeared to be important to the dolphin on four or more sets: the echo spectrum shape, the pattern of changes in target strength and number of highlights as a function of object orientation, and peak and center frequency. These data suggest that dolphins use multiple features and integrate information across echoes from a range of object orientations. *© 2006 Acoustical Society of America.*
[DOI: 10.1121/1.2161434]

PACS number(s): 43.80.Ka, 43.80.Lb [FD] Pages: 1867–1879

I. INTRODUCTION

Dolphins echolocate by emitting short, high-intensity clicks and processing the echoes reflected from objects. The ability of dolphins to detect and discriminate among objects via echolocation is well documented (for a recent review, see Au, 2000). Although these capabilities have been characterized, it is still not clear which acoustic features of object echoes are information-bearing parameters for dolphins, i.e., which features convey object properties such as size, shape, and material to an echolocating dolphin.

In a number of previous studies, experimenters have presented echolocating dolphins with discrimination tasks and then examined the echoes from the objects used in those tasks with simulated dolphin sonar signals (e.g., Au and Martin, 1988; Au and Pawloski, 1992; Au and Turl, 1991; Hammer and Au, 1980; Nachtigall *et al.*, 1980). An inspection of the object echoes allowed the experimenters to speculate about the echo features the dolphins may have used to discriminate between the objects. For example, Hammer and Au (1980) investigated a dolphin's ability to discriminate between hollow cylinders of the same diameter that varied in material composition (aluminum, bronze, glass, and steel). The dolphin performed well, and after the task was completed the cylinder echoes were obtained using simulated dolphin signals. A visual inspection of the cylinder echoes showed that the targets each had different arrival times for the secondary echo "highlight" (local maximum in echo amplitude). The researchers suggested that the predominant cue used by the dolphin in discriminating among the cylinders was probably time separation pitch generated by the first and second highlights. Humans, when presented with a pair of correlated sound pulses, perceive a pitch equal to $1/T$, where T is the time separation between pulses (Small and McClellan, 1963; McClellan and Small, 1965).

Such studies have been useful in determining a number of echo features that are *available* to dolphins asked to discriminate among objects that vary in size, shape, material, or structure (e.g., highlight structure of the echo waveform, time separation pitch, target strength, and frequency shifts in the peaks and nulls of the echo spectrum). However, it is not clear that dolphins actually attend to and utilize these echo features to give them information about object characteristics.

[a)]Present address: Brown University, Dept. of Neuroscience, Box 1953, Providence, RI 02912. Electronic mail: Caroline_Delong@Brown.edu

One way to determine which echo features dolphins actually use during a discrimination task is to take a quantitative approach to analyzing the object echoes. Instead of a visual inspection of the object features, the between-object differences in each of these features can be measured. In addition, the dolphin's error patterns during the discrimination task (e.g., how often each object is confused with each other object) can be analyzed in conjunction with the between-object variation in acoustic features to identify the acoustic features that the dolphin may have used. Two objects are confused to the degree that they share similar features. The combination of error patterns and acoustic similarity patterns thus indicate the specific features that underlie the confusion. For example, imagine that object 1 and object 2 had echoes that were very close in amplitude, but objects 1 and 3 had echoes that were significantly different in amplitude. If the dolphin confused objects 1 and 2 but did not confuse objects 1 and 3, then it can be inferred that the dolphin may have used amplitude to make its decision. In contrast, if a dolphin did not confuse objects 1 and 2, but confused objects 1 and 3, then it suggests that the dolphin may have used a different acoustic feature to make its decision. This method of comparing the dolphin's error patterns with the acoustic similarities between the echoes was employed in the present study to investigate how the dolphin used the acoustic features of the echoes.

The approach of this study was to give a dolphin an echolocation matching task in which he was presented with a sample object and then had to choose the identical object from among three alternatives. The dolphin was presented with a variety of objects that differed along one or more dimensions (e.g., size, shape, material, texture). After the matching task was completed, the object echoes were recorded in a test tank by projecting a dolphin click at the objects and the acoustic features of the echoes were measured. This study explored two hypotheses regarding the way dolphins use acoustic information in echoes: use of a single feature (hypothesis no. 1), or use of a linear combination of multiple features (hypothesis no. 2). Since there is no one-to-one correspondence between object characteristics (e.g., size, shape), and echo features (e.g., target strength, number of highlights), we predicted that the dolphin would use a combination of multiple acoustic features instead of a single acoustic feature.

II. BOTTLENOSE DOLPHIN ECHOIC MATCHING EXPERIMENT

A. Animal subject

The subject was an adult male Atlantic bottlenose dolphin (*Tursiops truncatus*) housed at Disney's Epcot's Living Seas in Orlando, Florida. At the beginning of the study the dolphin Toby was approximately 20 years old. Toby was an experienced research subject (Bauer and Johnson, 1994; Xitco, 1996) and had extensive experience with echoic matching (Xitco and Roitblat, 1996) and cross-modal matching (Harley *et al.*, 2003). Sessions were conducted in the main tank of the Living Seas (circular salt-water aquarium about 67 m in diameter and 9 m deep, with a volume of 22 million liters). During sessions, Toby usually received one-quarter of his daily allotment of approximately 9.5 kg of fish (herring [*Culpea harengus*], mackerel [*Scomber japonicus*], and capelin [*Mallotus villosus*]).

B. Materials and procedure

The stimuli were 27 assorted hardware-store objects that were unfamiliar to the dolphin (Fig. 1). The objects varied in size from 4.8 cm height $\times$ 4.8 cm width $\times$ 1.3 cm depth (smallest object: Small Stone Square) to 38.1 cm $\times$ 20.7 cm $\times$ 9.8 cm (largest object: Large Strainer). The objects were organized into nine object sets and each set contained three objects. Objects within each set were selected to vary along one nominal dimension: size, shape, material, or texture. Two object sets were selected for size differences (Strainers and Stone Squares), two object sets were selected for shape differences (Foam Cones and Stone Shapes), two object sets were selected for material differences (Rods and Figure 8's), and three object sets were selected for texture differences (Wooden Plaques, Green Foam, and Sockets). However, the objects were natural stimuli so they did not vary along only one dimension. For example, the Stone Shapes varied in both shape and size.

Samples were presented in the underwater sample apparatus, a 0.7 m square PVC-framed box tightly covered in 10 mm black polyethylene that was acoustically transparent but visually opaque. Comparison stimuli were presented in the choice apparatus, a larger (front: 2.76 m by 0.90 m, sides: 0.98 m by 0.90 m) but similarly fashioned rectangular structure, the top of which was positioned one meter below the water's surface (also covered in black polyethylene). The three comparison stimuli were centered within three square sections in the rectangular structure.

The basic procedure involved a three-alternative identity match-to-sample task in which the dolphin was presented with a sample stimulus, the sample was removed, and then the dolphin was required to select the identical stimulus from among three comparison stimuli. The objects were only accessible to the dolphin using echolocation. At the beginning of each trial, the dolphin investigated the sample as long as he wished at the underwater sample apparatus, after which he swam to the choice array located several meters behind him. When inspecting both the sample and the choices, the dolphin's head was unrestrained and he was free to ensonify the objects from several orientations (although he was not allowed to swim behind the sample or choice apparatus). After the dolphin positioned himself in front of his object of choice for about 3 s an assistant naïve to the sample's identity reported the dolphin's choice to the trainer who blew a whistle and reinforced the dolphin with two small fish for a correct choice, or tapped on a metal platform to recall him for an incorrect choice. Intertrial intervals averaged approximately 60 s (minimum 30 s).

Object sets were presented in the following order: Strainers, Wood Plaques, Foam Cones, Figure 8's, Stone Squares, Stone Shapes, Sockets, Green Foam, and Rods. Two 18-trial sessions were presented for each object set. Within a session, each stimulus was presented as the sample object an

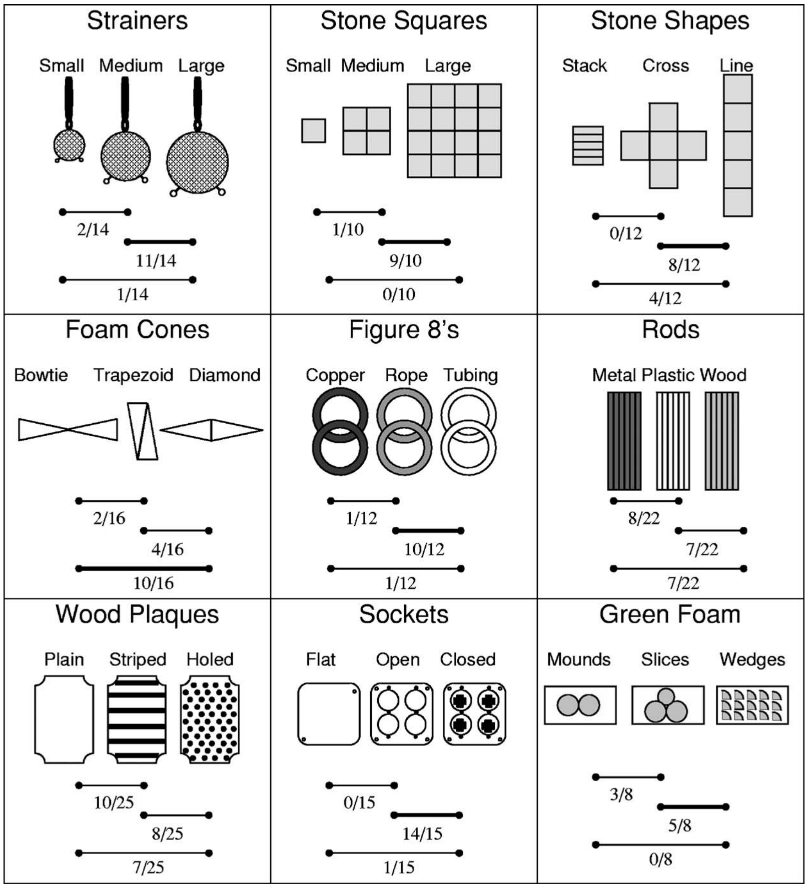

FIG. 1. Object sets. The dolphin's errors are shown for each object pair by lines connecting those objects. The number of errors the dolphin made for each pair is given under the line. The bold line indicates the "predominant" error (see text for details).

equal number of times. Each object was also presented equally often in each choice position. The order of the trials was randomized. The dolphin's ability to discriminate among the objects was measured using both percent correct matches and the unbiased sensitivity parameter, d' (Green and Swets, 1988).[1] The use of d' is preferable to percent correct matches because discrimination is inferred from nonrandom responding given the comparison choices, but does not require the subject to select the matching choice alternative. In other words, the dolphin could still show evidence of being able to discriminate the objects, even if it employs a perverse response rule (e.g., choosing object 2 every time object 1 is presented).

C. Results and discussion

Overall choice accuracy for all nine sets is as follows: Strainers (61%), Stone Squares (72%), Stone Shapes (67%), Foam Cones (56%), Figure 8's (67%), Rods (39%), Wood Plaques (31%), Sockets (58%), and Green Foam (78%). Chance choice accuracy is 33%, because the dolphin could choose from among three alternatives. The alpha level used to determine significance throughout this article was 0.05. A binomial test rather than parametric statistics was used because of the single subject design. The dolphin's choice accuracy was significantly above chance for all object sets except for the Rods and Wood Plaques (summed binomial test).

Figure 1 displays the dolphin's errors for each object set (under the lines connecting each pair of objects). In a match-to-sample task, errors occur when the dolphin chooses an object that does not match the sample object. For example, when the sample is "Large" and the dolphin chooses "Medium" the dolphin has made a Large-Medium error (Large-Medium and Medium-Large errors are collapsed into one group and shown under the line connecting those objects in Fig. 1). Chi-square tests were used to determine whether the dolphin's errors were distributed uniformly among the three possible object confusions. In the seven object sets in which the dolphin's overall choice accuracy was above chance, it confused one of the three object pairs more frequently than the other two pairs (Stone Squares $[\chi^2(2)=14.6, N=10, p<0.001]$, Strainers $[\chi^2(2)=13.0, N=14, p<0.01]$, Stone Shapes $[\chi^2(2)=8.0, N=12, p<0.05]$, Foam Cones $[\chi^2(2)=6.5, N=16, p<0.05]$, Figure 8's $[\chi^2(2)=13.5, N=12, p<0.01]$, Sockets $[\chi^2(2)=24.4, N=15, p<0.001]$, Green Foam $[\chi^2(2)=6.8, N=8, p<0.05]$).

The pair of objects the dolphin confused most often was called the predominant error (defined as the error type that accounts for more than 50% of the errors, and shown as bold lines in Fig. 1). For example, in the Stone Squares set, the dolphin's predominant error was to confuse the Medium and Large objects. In the other two object sets, Rods and Wood Plaques (overall choice accuracy below chance), the dolphin's errors were distributed uniformly among the three possible object confusions (Rods $[\chi^2(2)=0.1, N=22, p>0.05]$, Wood Plaques $[\chi^2(2)=0.6, N=25, p>0.05]$). For these sets, there was no predominant error because the dolphin confused all three pairs of objects.

III. ACOUSTIC MEASUREMENTS OF OBJECTS

A. Materials and procedure

After the completion of the cross-modal matching experiment with the dolphin, echoes from all the objects were obtained using a representative click recorded from a different male bottlenose dolphin. This dolphin click, which has been used in numerous studies (see Au, 1993), was 70 μs long with a peak frequency of about 120 kHz and a 60-kHz bandwidth. Dolphin clicks can vary in frequency and intensity, but this click was considered to be a typical, average click that was likely to be very similar to the clicks that were used by the dolphin Toby during this study (Au, 1993). The object echoes were obtained using this recorded dolphin click because recording real-time echoes from the dolphin's own clicks during the matching task was not feasible given the time constraints of this study (see general discussion).

The measurements took place in a seawater tank (with the exception of the extremely buoyant Foam Cones which were measured in an open ocean pen but otherwise in a similar fashion; see Fig. 1 in Benoit-Bird *et al.*, 2003). The cylindrical tank was 1.82 m in height by 2.41 m in diameter and contained 8.3 m^3 of seawater. The transmitting transducer and the receiving hydrophone were mounted on the same transducer assembly (a 28.55 cm $\times$ 20.30 cm aluminum plate). The transmitting transducer was located 2.54 cm above the receiver. The transducer and receiver were custom built with piezoelectric ceramic circular disks that were 6.35 cm in diameter and 6.35 mm thick (Material Systems).

The dolphin click signal was generated by a Quatech WSB-10 function generator board housed in a PC and amplified (Hafler P3000 Transnova). The received echoes were gated, amplified using a custom-built amplifier, and filtered before being digitized at 1 MHz using a Rapid Systems R1200 A/D Converter. An oscilloscope was used to view the signals during echo collection (Tektronix TDS 210).

The transducer assembly and the targets were placed 1 meter below the water's surface. Individual targets were hung with monofilament line from a T-bar suspended over the surface of the water. The T-bar was linked to a calibrated rotor (ILC Data Device Corp. API 30602) that could rotate the targets 360 degrees in 1.3 degree increments. Some of the objects required lead weights (343 or 571 g) attached with monofilament line to make them negatively buoyant. These weights were hung 62.23 cm below the targets to minimize their contribution to the echoes. These weights were not hung from the targets during the dolphin echoic matching task. Weights were not needed during the echoic matching task because the targets were held in place in the target boxes using thin polyethylene strings attached to the top and bottom of the targets (not possible in the measurement tank).

In order to simulate the experience the dolphin could get from swimming by the target and ensonifying it from different angles, echoes were measured from multiple angles parallel to the horizontal axis of the target. Multiple aspect angles were measured because the dolphin could ensonify the object from different angles relative to the object during the matching task and because there is evidence that dolphins attend to the pattern of changes in acoustic features as an object is scanned across a range of target orientations (e.g., Nachtigall *et al.*, 1980).

It was estimated that the dolphin was exposed to a range of angles spanning +/− 15 degrees (30 degree span) during the matching task (videotapes of some sessions suggest that the dolphin did not consistently investigate a wider range of angles). Simulating this range, each measurement of an object produced a 23-echo train in which one echo was captured for each angle (1.3 degrees apart) between −15 degrees and +15 degrees (the orientation of the target that faced the front of the choice apparatus in the matching task was called the 0 degree angle). Ten echo train measurements were collected for each object (23 echoes per train $\times$ 10 trains per object $\times$ 27 objects = 6210 total echoes). All ten measurements for a single object were completed consecutively on the same day (all ten measurements were made for one object before measuring another object). The objects were measured in random order.

1. Acoustic feature extraction

Six acoustic features were extracted from each individual echo: (1) target strength, (2) number of highlights, (3) duration, (4) peak frequency, (5) center frequency, and (6) rms bandwidth (Figs. 2(a) and 2(b)). These features characterize the echo in the time and frequency domains and are standard features that have been examined in other studies (for a review see Au, 1993).

Target strength is defined as the ratio in dB of the echo intensity measured 1 m from the target to the intensity of the incident signal at the location of the target. Highlights, defined as local amplitude maxima in the time domain, were localized and counted by finding the envelope of the signal, smoothing the envelope (by averaging 15 points centered around the sample), thresholding the smoothed envelope, taking the derivative of the thresholded envelope, representing the derivative as a trinary signal (+1, 0, −1), taking the derivative of the trinary signal, and extracting zero crossings that indicate slope changes associated with highlight onset and peak events. The number of highlights was simply defined as the number of peaks detected.

Echo duration was calculated by finding the envelope of the signal, thresholding the envelope at 1 standard deviation, then finding the time between the first and last suprathreshold sample. The peak frequency is the frequency of the signal at which the spectrum has its maximum value. The center frequency is defined as the frequency that divides the power spectrum into two equal energy parts. The root mean square (rms) bandwidth indicates the frequency range around the center frequency in which the majority of spectral energy lies.

Since there is evidence that dolphins attend to the pattern of changes in acoustic features as an object is scanned across a range of orientations (Nachtigall *et al.*, 1980), additional acoustic features were extracted from echo trains instead of individual echoes (one train is 23 echoes collected from −15 to +15 degrees; Figs. 2(c) and 2(d)). Four features were calculated: relative target strength, relative number of highlights, target strength "bumpiness," and highlight

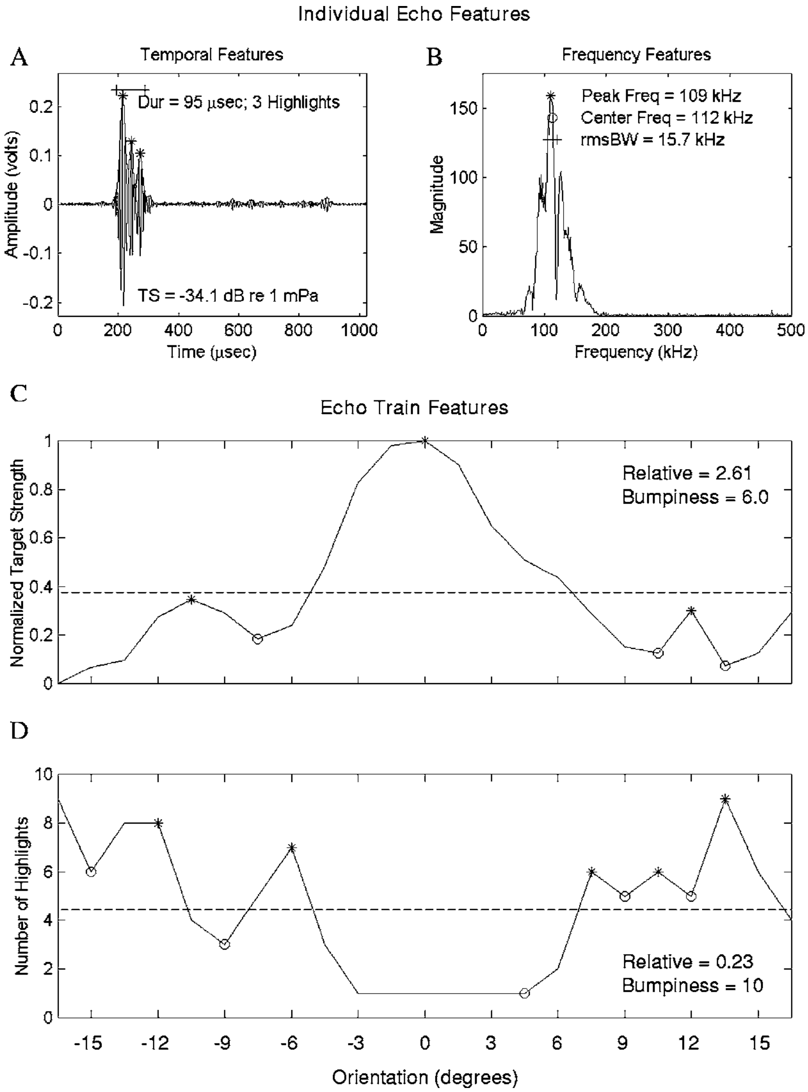

FIG. 2. Features that were extracted from the echoes. Parts A and B show features extracted from individual echoes. This example echo is from the Cross. Part A shows the echo in the time domain and the two temporal features that were extracted: duration (line above echo) and number of highlights (marked with asterisks). Target strength is also shown on the bottom of the graph. Part B shows the echo in the frequency domain and the three frequency features that were extracted: peak frequency, center frequency, and rms bandwidth. Parts C and D show features extracted from echo trains. These example data are from the Flat Socket. Part C shows the normalized target strength plotted as a function of orientation. Part D shows number of highlights plotted as a function of orientation. The dotted line shows the average target strength (or number of highlights) from −15° to +15°. The relative target strength (or number of highlights) score is calculated by dividing the value at the 0° orientation by the value averaged from −15° to +15°. The stars and open circles show changes in line slope. The total number of changes in slope across −15° to +15° yields the "bumpiness" score.

"bumpiness." These features were extracted from each of the ten echo trains for each object. The relative target strength (or number of highlights) score is calculated by dividing the value at the 0 degree orientation by the average value for all orientations. These two features provide a measure of target strength or highlight variability relative to the 0 degree aspect. Target strength (or highlight) bumpiness was defined as the number of slope direction changes across all orientations (see Figs. 2(c) and 2(d)). These echo train features have not been systematically examined in other studies.

B. Results and discussion

1. Variation in single acoustic features of the echoes

a. Changes in acoustic features of the echoes within objects as a function of orientation. Figure 3 shows changes in the echo structure in the time domain as a function of orientation for one of the object sets—the Stone Squares. The way the echo changed as a function of orientation depended on object type. Since echo highlights are caused by reflections from different parts of the target, more complexly shaped targets produce echoes with more complex structures. Whether the waveforms were short and simple or long and complex depended on the number of reflective surfaces presented by the object at each particular angle.

To examine whether acoustic features changed significantly as a function of orientation, separate multivariate analyses of variance (MANOVA) were performed for each object with one independent variable (orientation) and six dependent variables (target strength, number of highlights, peak frequency, center frequency, bandwidth, and duration). A significant multivariate effect of orientation was found for all 27 objects (for all objects, $df=132$, 1242, $p<0.001$). All six acoustic features varied significantly as a function of orientation for all 27 objects (for all objects, $df=22$, 230, $p<0.001$). Each of the acoustic features changed in different ways as a function of orientation, and the patterns of change also varied between objects. An examination of all the object sets revealed that there were many different ways the acoustic features could change as a function of orientation. For example, in the Stone Squares set, target strength increased as the objects rotated toward 0 degrees, whereas the number of highlights and duration either decreased at 0 degrees or remained constant across orientations (see Fig. 4).

b. Acoustic differences between the objects. To examine between-object differences in the acoustic features, separate

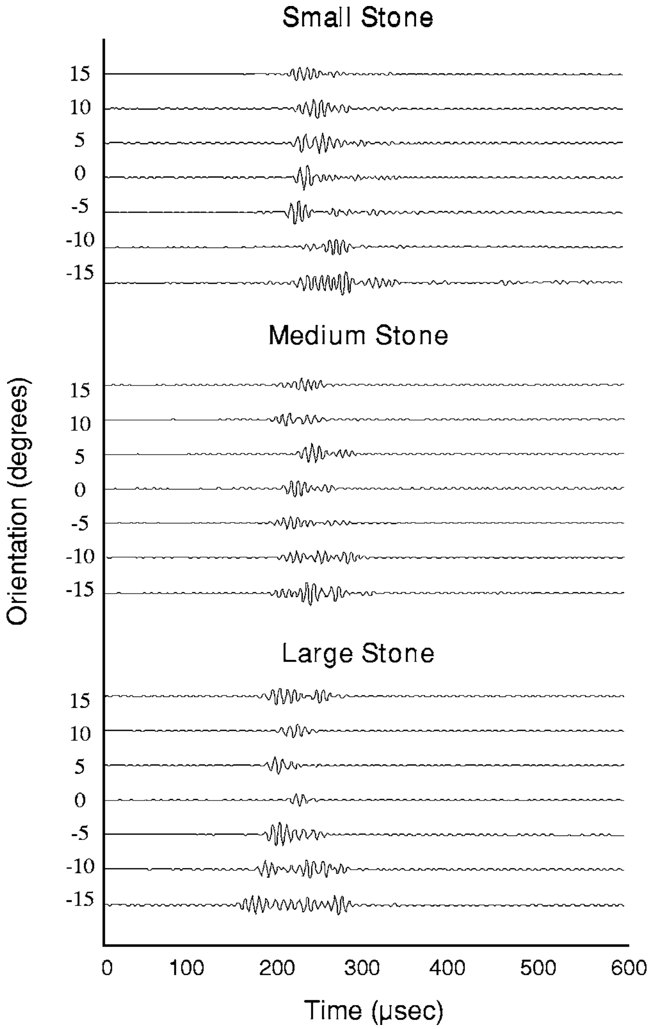

FIG. 3. Echo waveforms for the Stone Squares set showing changes in the echoes of each object as a function of orientation. The orientation facing the front of the choice apparatus during the echoic matching task is called 0°. Amplitude was normalized so amplitude differences between the three objects cannot be seen.

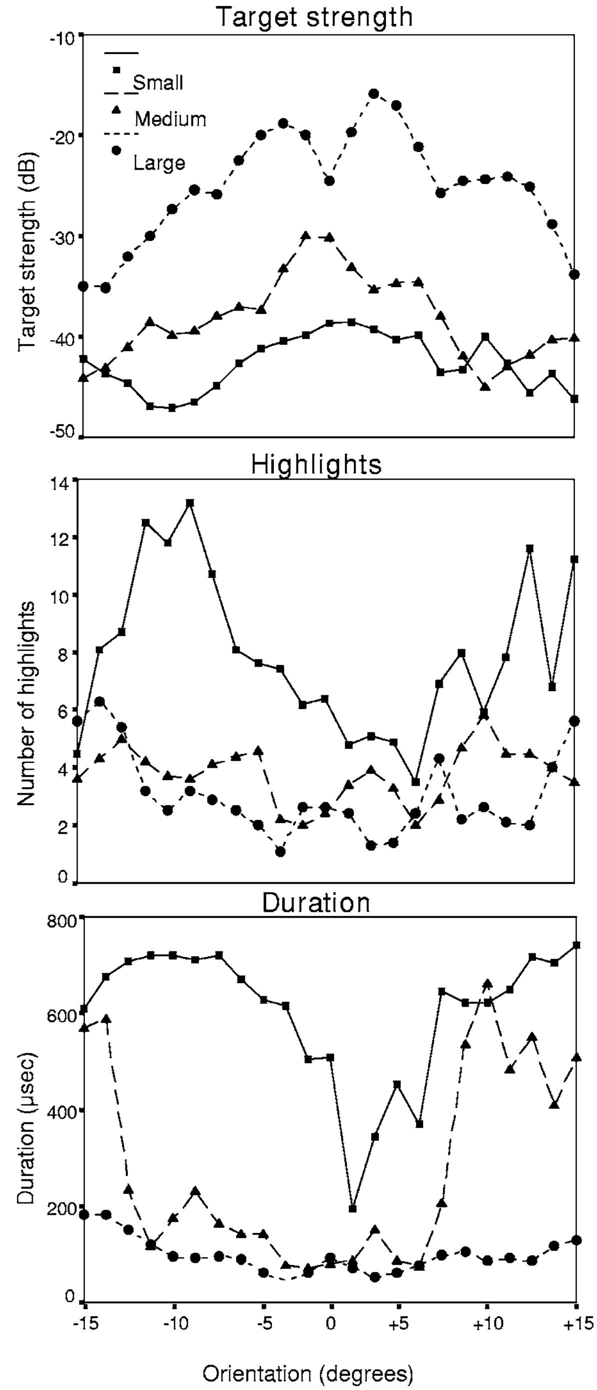

FIG. 4. Changes in three echo features (target strength, number of highlights, duration) as a function of object orientation for the Stone Squares set. The orientation facing the front of the choice apparatus during the echoic matching task is called 0°.

multivariate analyses of variance (MANOVA) were conducted for each of the nine object sets. Between-object differences were examined within object sets (instead of examining differences between all possible pairings of the 27 total objects) because the dolphin was asked to match objects within sets but not between sets. Two separate groups of MANOVAs were conducted for (1) the six acoustic features of individual echoes and (2) the four acoustic features of echo trains. The features of the individual echoes were not analyzed together with the features of the echo trains because the number of samples was different for individual echoes versus echo trains (230 samples versus 10 samples). For the individual echo analysis, acoustic features were extracted from 230 individual echoes per object, one echo from each angle in the range of −15 to +15 degrees (23 echoes) $\times$10 trains. For the echo train analysis, acoustic features were extracted from 10 echo trains per object (i.e., 10 samples per object).

In the echo MANOVAs, the independent variables were object and orientation and the dependent variables were the six acoustic features of individual echoes. There were significant multivariate effects of object, orientation, and object$\times$orientation for all nine object sets (for all tests, $p < 0.001$). The acoustic features of the objects differed among the different objects, and the differences among objects changed as the orientation of those objects changed. For seven of the nine object sets (Stone Squares, Strainers, Stone Shapes, Foam Cones, Rods, Sockets, Green Foam), there were significant differences among objects for all six acoustic features (for all features in all sets, $p < 0.001$). For the Figure 8's set, the objects were not significantly different in

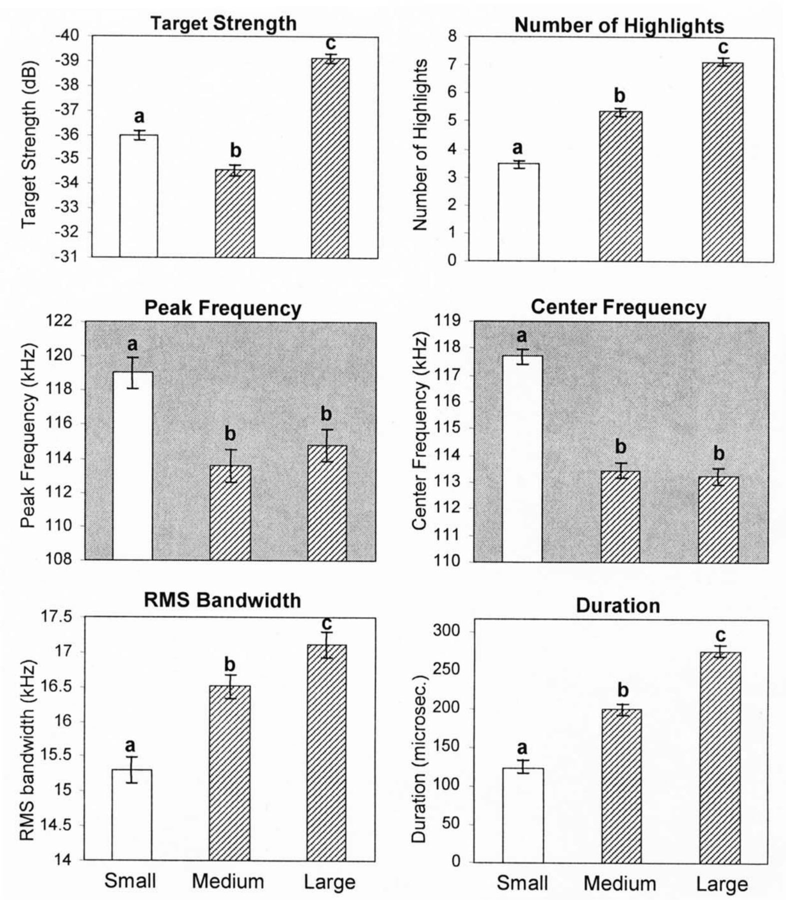

FIG. 5. Features of individual echoes for the Strainers set. The bars representing the objects that the dolphin confused most often (predominant error) are cross-hatched (Medium and Large). Objects sharing the same letter were not statistically significantly different from each other ($p<0.05$). Panels with features that can account for the dolphin's errors are shaded gray.

bandwidth, but the objects were significantly different across all other features ($p<0.001$). For the Wood Plaques set, the objects were not significantly different in peak frequency, but the objects were significantly different across all other features ($p<0.001$).

To examine which pairs of objects within each set showed significant differences in these acoustic features, post-hoc object comparisons were conducted. Figure 5 displays the results of the post-hoc analyses for one of the nine object sets: the Strainers. The height of the vertical bars represents the values for each feature averaged across all object orientations. In Figs. 5 and 6, statistically significant differences between means are indicated by different lower case letters. For example, in the Strainers set, the Small Strainer, Medium Strainer, and Large Strainer are all significantly different in target strength, number of highlights, bandwidth, and duration. The Medium Strainer and the Large Strainer are not significantly different in peak frequency and center frequency, but both are different from the Small Strainer.

To examine which acoustic features the dolphin may have used to discriminate among the objects, the errors made by the dolphin were considered in conjunction with the results of the post-hoc object comparisons. In Fig. 5, the bars representing the objects that the dolphin confused most often (predominant errors) are cross-hatched. If two objects in a set had similar values for a certain acoustic feature (i.e., they were *not* significantly different on that feature), *and* the dolphin made errors between those two objects, then it was inferred that those acoustic features may have been part of the dolphin's decision making process and used by the dolphin to construct representations of the objects. In Fig. 5, graphs with acoustic features inferred to be part of the dolphin's decision making process according to the above logic are shaded in gray. For example, in the Strainers set, the dolphin confused the Medium and Large Strainer, which are similar in peak and center frequency. Thus, peak and center frequency may have been important acoustic features for the dolphin for that object set.

A second group of MANOVAs were conducted to investigate between-object differences in acoustic features calculated from echo trains. A separate MANOVA was conducted for each object set with one independent variable (object) and four dependent variables (relative target strength, relative number of highlights, target strength bumpiness, highlight bumpiness). A significant multivariate object effect was found for all nine object sets (for all sets $df=8$, 50, $p<0.001$). For four of the nine object sets (Strainers, Figure 8's, Rods, Green Foam), there were significant differences among objects for all four acoustic features ($df=2$, 27, $p<0.01$). For three object sets (Stone Squares, Foam Cones, Sockets), the objects were not significantly different in highlight bumpiness (the three other features were significantly different, $p<0.01$). For the final two sets (Stone Shapes, Wood Plaques), the objects were not significantly different in target strength bumpiness (the three other features were significantly different, $p<0.001$). To examine which pairs of objects within each set showed significant differences in these acoustic features, again post-hoc object comparisons

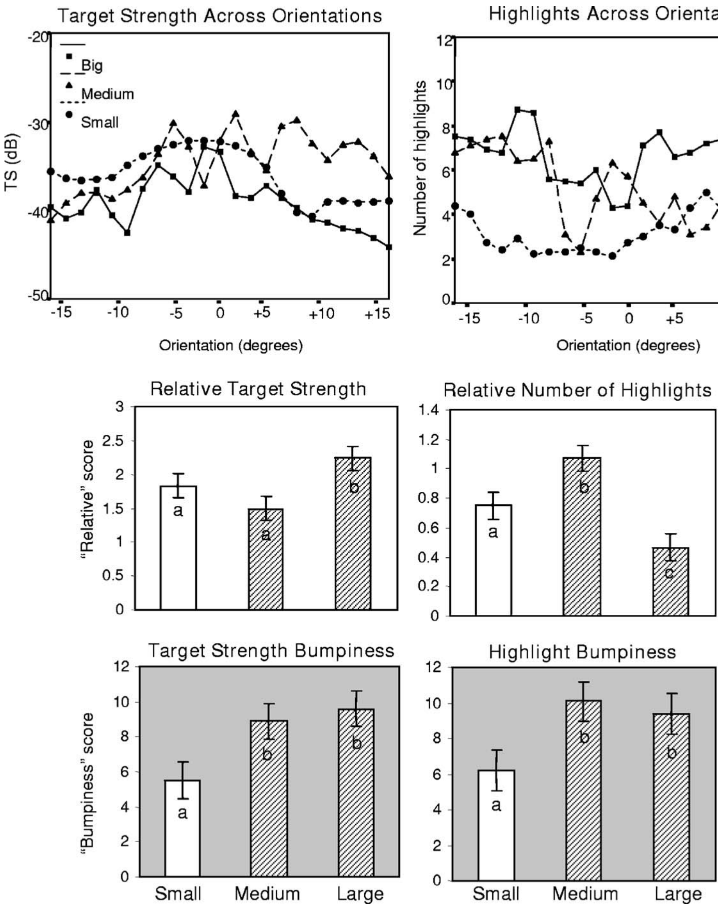

FIG. 6. The top two graphs show the change in target strength and number of highlights as a function of orientation for the Strainers set. The bottom four graphs show four features extracted from the echo trains. The bars representing the objects that the dolphin confused most often (predominant error) are cross-hatched (Medium and Large). Objects sharing the same letter were not statistically significantly different from each other ($p<0.05$). Panels with features that can account for the dolphin's errors are shaded gray.

were conducted and considered in conjunction with the dolphin's errors. Figure 6 shows the results for the Strainers set.

Table I summarizes the features that may explain the dolphin's errors. Note that different acoustic feature(s) were implicated to be important to the dolphin for each object set. There was not a single acoustic feature or two that seemed to be used by the dolphin for every object set. These results suggest that the dolphin did not use a single acoustic feature to discriminate among all types of objects (hypothesis no. 1). Four features appeared to be important to the dolphin on four or more object sets: target strength bumpiness, highlight bumpiness, peak frequency, and center frequency. In eight of the nine object sets, the dolphin's error patterns implied the use of multiple acoustic features.

c. Comparison of variance within objects to variance between objects. The variability in acoustic features within objects (as a function of orientation) was compared to the variability in acoustic features between objects. The variance

TABLE I. Acoustic features that may explain the dolphin's errors for each of the object sets.

	Object set									
Acoustic features	Stone squares	Strainers	Stone shapes	Foam cones	Figure 8's	Rods	Wood	Sockets	Green foam	Total
Individual Echo										
Target strength			X	X						2
Highlights (#)				X						1
Peak frequency		X		X	X	X	X		X	6
Center frequency		X		X			X	X	X	5
RMS bandwidth			X	X						2
Duration				X						1
Echo Train										
Relative TS			X						X	2
Relative HL					X	X		X		3
TS Bumpiness		X		X		X	X		X	5
HL Bumpiness		X			X	X	X			4
Total	0	4	3	7	3	4	4	2	4	
Dolphin's Accuracy	72%	61%	67%	56%	67%	39%	31%	58%	78%	

was computed for each three-object set ("between-object variance") and then for each object within each set separately ("within-object variance"), with each echo as a single observation. A ratio of within:between object variance was calculated by dividing within-object variance by between-object variance. If the value of this ratio was greater than 1.0, then the within-object variance was greater than the between-object variance.

For one of the object sets in which the dolphin performed significantly above chance (Stone Shapes: 67% correct), the within-object variance was greater than or equal to the between-object variance (ratios ranged from 1.30 to 0.90) for four features for Cross (target strength, number of highlights, center frequency, bandwidth), five features for Line (target strength, peak and center frequency, bandwidth, duration), and three features for Stack (number of highlights, peak frequency, duration). These results show the dolphin succeeded at the matching task even when the acoustic variance within objects (as a function of orientation) was greater than or equal to the variance between objects.

2. Spectral correlations

To test the hypothesis that the shape of the echo spectrum was a feature used by the dolphin, the dolphin's errors were examined alongside similarities between the entire echo spectra of objects within each set. If the dolphin were using the shape of the spectrum, it would be expected to confuse objects with similar spectra. The echo spectra from object echoes at the 0 degree aspect (the aspect of the object that faced the front of the echoic choice apparatus) were normalized such that each individual spectrum had a maximum power of 1.0, and averaged over 10 measurement runs per object. These average spectra for each object were correlated with each other object within an object set. Object spectra were not correlated with objects outside their object set because the dolphin was asked to match objects within sets but not between sets. If the dolphin was using the shape of the spectra, we would expect the two objects that were the most often confused to have the highest *r* value. The *r* values for the correlations are shown in Table II alongside the proportion of errors made by the dolphin in each error type.

For four of the nine object sets, the two objects with the highest correlation between echo spectra were the two objects most often confused by the dolphin. For example, in the Green Foam set, the dolphin's predominant errors were between Slices and Wedges (63% of the errors). Slices and Wedges were more highly correlated ($r=0.952$) than either Slices and Mounds ($r=0.670$) or Mounds and Wedges ($r=0.635$). The Wood Plaques are considered to be one of these four sets because all three objects are highly correlated and all three objects are confused by the dolphin. For the other five object sets, the two objects with the highest correlation between echo spectra were *not* the two objects most often confused by the dolphin. For example, in the Stone Squares set, the dolphin's predominant errors were between Medium and Large (90% of the errors). However, Small and Large were more highly correlated ($r=0.787$) than Medium and Large ($r=0.425$) or Small and Medium ($r=0.365$).

TABLE II. Spectral correlations and the dolphin's errors.

Object set	Error type	Correlations *r* values	Errors Proportions
Strainers	Small-Medium	0.570	0.142
	Medium-Large	**0.708**	**0.786**
	Small-Large	0.580	0.071
Stone squares	Small-Medium	0.365	0.100
	Medium-Large	0.425	**0.900**
	Small-Large	**0.787**	0
Stone shapes	Cross-Line	0.858	**0.667**
	Line-Stack	0.856	0.333
	Cross-Stack	**0.952**	0
Foam cones	Diamond-Bow	**0.992**	**0.625**
	Trap-Diamond	0.914	0.250
	Trap-Bow	0.920	0.125
Figure 8's	Rope-Tubing	0.722	**0.833**
	Copper-Rope	0.698	0.083
	Copper-Tubing	**0.766**	0.083
Rods	Metal-Plastic	0.439	**0.364**
	Plastic-Wood	0.811	**0.318**
	Metal-Wood	**0.858**	**0.318**
Wood plaques	Holed-Plain	**0.995**	**0.280**
	Striped-Plain	**0.997**	**0.400**
	Hole-Striped	**0.999**	**0.320**
Sockets	Closed-Open	0.838	**0.933**
	Closed-Flat	**0.948**	0.067
	Open-Flat	0.906	0
Green foam	Slices-Mounds	0.670	0.375
	Slices-Wedges	**0.952**	**0.625**
	Mounds-Wedges	0.635	0

Note: The dolphin's predominant errors are highlighted in bold. When the dolphin did not make one type of error predominantly (Rods, Wood), all errors are highlighted. The object pairs with the highest correlation are also highlighted in bold. Object sets in which the highest correlation matches the predominant confusion are shown in bold and italics.

3. Linear model of the dolphin's performance

To test the hypothesis that the dolphin used a linear combination of multiple acoustic features, the dolphin's<> choice accuracy was compared to the classification accuracy of a statistical pattern classification model: the discriminant function analysis (DFA). In a DFA, one or more continuous predictor variables are used to form a linear model, from which stimulus classification is made. A separate DFA was run for each object set. The DFA was used to predict object type given the six single-echo acoustic features. For each object within a set, there were 230 echoes (23 object orientations $\times$ 10 measurement runs).

To compare the DFA results with the dolphin's classification results (12 trials per object in each set), the classification percentages for the original DFA counts were multiplied by 12. The classification matrices for the dolphin are presented above the transformed classification matrices for the DFA for three of the object sets in Table III. The DFA models were more accurate than the dolphin in overall matching accuracy, with the exception of the Green Foam object set. The classification performance of the DFA was compared against the dolphin's performance using chi square tests with the dolphin's choices as the expected distribution. The pattern

TABLE III. Classification matrices for the dolphin vs the DFA.

		Choice		
STONE SQUARES				
Dolphin's Performance		Small	Medium	Large
	Small	12	0	0
Sample	Medium	1	2	9
	Large	0	0	12
DFA Echo Classification		Small	Medium	Large
	Small	10.2	1.8	0.0
Sample	Medium	1.9	10.1	0.0
	Large	0.4	1.0	10.6
FIGURE 8'S				
Dolphin's Performance		Copper	Rope	Tubing
	Copper	12	0	0
Sample	Rope	1	5	6
	Tubing	1	4	7
DFA Echo Classification		Copper	Rope	Tubing
	Copper	9.9	0.0	2.1
Sample	Rope	0.0	11.5	0.5
	Tubing	0.0	1.4	10.6
SOCKETS				
Dolphin's Performance		Closed	Flat	Open
	Closed	6	0	6
Sample	Flat	1	11	0
	Open	8	0	4
DFA Echo Classification		Closed	Flat	Open
	Closed	10.7	0.7	0.6
Sample	Flat	0.9	10.8	0.3
	Open	1.5	0.9	9.5

of classification made by the DFA was significantly different from the dolphin for all of the sets except the Stone Shapes, Foam Cones, and the Green Foam (Stone Squares $[\chi^2(4)=42.86, N=36, p<0.001]$, Strainers $[\chi^2(4)=68.73, N=36, p<0.001]$, Stone Shapes $[\chi^2(4)=6.34, N=36, p>0.05]$, Foam Cones $[\chi^2(4)=3.87, N=36, p>0.05]$, Figure 8's $[\chi^2(4)=19.53, N=36, p<0.001]$, Rods $[\chi^2(4)=31.26, N=36, p<0.001]$, Wood Plaques $[\chi^2(4)=21.62, N=36, p<0.001]$, Sockets $[\chi^2(4)=21.63, N=36, p<0.001]$, Green Foam $[\chi^2(4)=2.98, N=36, p>0.05]$).

The DFA models did not produce the same biased classification patterns made by the dolphin for several object sets. For example, in both the Stone Squares set and the Strainers set, when the dolphin was presented with the Medium object he consistently chose the Large object. The DFA model did not show the same choice bias with the Medium object, it correctly classified the Medium object on most trials. For other object sets, the DFA did not replicate the pattern of errors made by the dolphin. For example, in the Figure 8's set, the dolphin's errors consisted primarily of confusions between Rope and Tubing (10/12 errors). In contrast, the DFA rarely misclassified Rope and Tubing. Another example is the Sockets set. The majority of the dolphin's errors were confusions between the Open Socket and the Closed Socket (14/15 errors). The DFA rarely misclassified the Open and Closed sockets.

When the classification performance of the DFA and the dolphin's performance were compared using values of d', the results show that the performance of the DFA and the dolphin were not similar. For two of the object sets (Rods and Wood), the DFA was more accurate in matching all of the object pairs (DFA's d' values higher than the dolphin's). For two of the object sets (Stone Shapes and Green Foam), the DFA was less accurate than the dolphin in matching all of the object pairs. The DFA's d' value was subtracted from the dolphin's d' value for each of the 27 object pairings (3 pairs $\times$ 9 object sets) to yield a difference score (e.g., for Small Stone vs Medium Stone: dolphin $d'=3.53$ vs DFA $d'=2.03$ for a difference score of 1.5). When the difference scores were between -0.3 and 0.3, the d' values were considered to be similar (scores ranged from -3.34 to 4.98). Out of the 27 object pairings, the d' values for the DFA were similar to the d' values for the dolphin for only 4 object pairings [Bow-Diamond (-0.2), Diamond-Trapezoid (0.29), Copper-Tubing (0.21), Metal-Plastic (-0.25)].

IV. GENERAL DISCUSSION

The present study explored two hypotheses regarding the way dolphins use acoustic information in echoes: (1) use of a single feature, or (2) use of a linear combination of multiple features. The results suggested that dolphins do not use a single feature across all object sets or a linear combination of the individual echo features. Five features appeared to be important to the dolphin on four or more object sets: the shape of the echo spectrum, the pattern of changes in target strength and number of highlights as a function of object orientation, and peak and center frequency. In eight of the nine object sets, the dolphin's error patterns implied the use of multiple acoustic features. Thus far, researchers have assumed that dolphins made use of multiple features in their echoes (Au, 1993), and these data provide some evidence for that assumption.

The first hypothesis included two possibilities: (A) use of a single, stable acoustic feature in all circumstances, and (B) use of a single acoustic feature, but the feature shifts based on the characteristics of the objects. For example, in case A, the dolphin always uses target strength regardless of whether objects vary in size, shape, or material, and in case B the dolphin uses peak frequency to discriminate between objects that vary in material and target strength to discriminate between objects that vary in size. The results of this study show that the dolphin did not use a single acoustic feature for all object sets, which eliminates case A. These data cannot completely rule out case B: the possibility that a dolphin may use a single, shiftable feature to discriminate among the objects (see Table I). However, the Stone Square set results argue against the use of a single feature since there was no single feature that could explain the dolphin's errors across all 12 measured features. The single feature hypothesis is improbable because there is no simple one-to-one relationship between the physical features of objects (e.g., size, shape, material) and the acoustic features of echoes (e.g., amplitude, frequency). For example, echo amplitude does not map only to the size of the object, it is affected by the both the size and material of the object (along with other characteristics).

It would be more advantageous for a dolphin to use multiple acoustic features in the echoes than to use a single acoustic feature. One big advantage to using multiple features is that the values of some features would disambiguate other features. For example, if two objects had very similar target strengths, but different highlight structures, a dolphin attending to more than one feature would be able to discriminate between those objects. In many of the object sets analyzed in this study, two objects had similar values for one or more acoustic features yet the dolphin was able to discriminate between them.

Using certain acoustic features to disambiguate other features becomes particularly important when considering the variation in features as a function of orientation. This variation in features as a function of orientation was at times greater than the between-object variation! A major finding in this study was that *all* the measured acoustic features for nearly all the objects varied as a function of orientation. This means the echoes from an object in one orientation are quite different from the echoes from the *same* object ensonified from a different orientation. Two objects could have the same value for an acoustic feature from one orientation, but different values of that same feature from another orientation. A dolphin has the best chance of picking up differences between objects if it attends to multiple features, and gets multiple "looks" at an object from different orientations (this would correspond to the dolphin moving its head or swimming by the object).

The fact that the acoustic features of objects can vary so much with orientation actually creates a potentially salient "feature" for the dolphin: the pattern of changes in the echo across several orientations of the object. In a previous study, it appeared that the dolphin attended to this feature to discriminate between objects that varied in shape (Nachtigall *et al.*, 1980). The present results indicate that the dolphin could have used the pattern of changes in the echo (target strength bumpiness and highlight bumpiness) in six of the nine object sets. Future studies should measure object echoes from various orientations and consider the strong possibility that the dolphin gathers and uses acoustic information from more than one orientation (as long as the dolphin's head is unrestrained as it was in this study).

The hypothesis that dolphins use multiple acoustic features was tested by comparing the dolphin's performance to a discriminant function analysis (DFA). The dolphin's classification patterns and accuracy did not match the DFA's, which suggests that the dolphin may not have used a linear combination of these particular features to discriminate among the objects. One interpretation of these data is that the dolphin does in fact use a linear combination of features, but the features extracted in this study were incorrect and/or insufficient. In the current study we extracted features from the time waveforms and the spectra. It is also possible to extract spectrogram features or other time/frequency derived features (e.g., Okimoto *et al.*, 1998).

Another possibility is that the dolphin used a combination of multiple features that was complex and nonlinear. It is very likely that dolphins combine echo features nonlinearly to recognize objects. Human cognition is considered to be nonlinear in many ways, particularly when it comes to pattern recognition and decision-making. Because artificial neural networks resemble the organization of biological neural systems, they have been used frequently to model the performance of biological systems. There have been several efforts to model the echolocation performance of dolphins using neural networks (e.g., Au *et al.*, 1995; Roitblat *et al.*, 1989). The networks in these studies used temporal and/or spectral features of the echoes and performed at times better or worse than the dolphin. The drawback of these studies is that they did not attempt to compare the error patterns of the dolphin and the network, so it is difficult to ascertain whether the network was using the same features as the dolphin. Since neural networks are capable of performing nonlinearly, this is a promising method of investigating how dolphins use acoustic features if the errors patterns of the dolphin and the network are analyzed.

It is a limitation of the current study that the sonar signals from the actual dolphin and the resulting echoes were not recorded and measured during the behavioral task. The acoustic features of the echoes are affected by not only the orientation of the objects, but also the characteristics of the dolphin's outgoing sonar signal. The objects in this study were measured in a test tank using a dolphin click that did not vary from trial to trial. In contrast, dolphin signals can vary in both frequency and intensity (Au, 1993). An individual dolphin's signals varied in peak frequency from 60 to 140 kHz (Au, 1980). Dolphins can vary the amplitude of their signals over a large dynamic range (e.g., peak-to-peak source levels vary from 150–230 dB re 1 μPa; Au, 1993). Dolphins modify their signals in an adaptive manner according to the environment and task demands. For example, dolphins will emit higher amplitude signals in noisy environments (Au *et al.*, 1974). In a situation where a dolphin is performing a specific task over and over again, there can be fluctuations in the characteristics of the emitted clicks during the task (Au, 1993). These variations in a dolphin's outgoing clicks will influence the acoustic characteristics of the resulting object echoes. For example, the amplitude of the echoes from an object is dependent on the source level of the projected click (Au, 1993).

To obtain the most accurate perspective on how dolphins use acoustic information in echoes, the actual outgoing signals and echoes should be recorded during the task. This was not done during the current study because it would have added additional animal training time that was not feasible given time constraints. To record the dolphin's signals in a free-swimming task such as this one, the dolphin must be trained to perform the matching task efficiently while carrying a bite-plate hydrophone. This difficult training procedure can take many months to a year. Because of these constraints, Toby's click signals were not recorded. It is assumed that he was using typical clicks resembling the recorded click that was used to obtain echoes for these acoustic measurements.

In the present study, statistical methods were used to determine between-object differences in echo features. A potential problem with this method is that statistical significance may not equal biological significance. Dolphins' ability to discriminate between sounds with different intensities is known to be approximately 1 dB (Vel'min, Titov, and Yurkevich, 1975, cited in Au, 1993). In the Figure 8's set, Rope has a target strength of −32.9 dB, which is *statistically* significantly different than the target strength for Tubing (−32.3 dB). However, dolphins may not perceive objects with target strengths within 1 dB of each other to be different since 1 dB is about the limit of their intensity discrimination ability. The Figure 8's set is the only object set in which objects that are separated by less than 1 dB are statistically significantly different. In all other sets, objects that are statistically different should also be perceived as different to the dolphin (i.e., objects are separated by 1 dB or more).

Dolphins' ability to discriminate between differences in the duration, frequency, bandwidth, or number of highlights in short broadband echoes is unknown. This makes it difficult to compare statistical and biological significance for these features. For example, the peak frequency of the Small Strainer (119.1 kHz) was statistically significantly different than the Medium (113.6 kHz) and Large Strainer (114.6 kHz). Since dolphins' frequency discrimination ability for click signals has not been comprehensively assessed, it is unknown whether the dolphin could actually discriminate between an echo with a peak frequency of 114 kHz and an echo with a peak frequency of 119 kHz. Dolphins can discriminate between a 114.0 kHz tone and a 114.2 kHz tone (Herman and Arbeit, 1972). Assuming that dolphins' frequency discrimination ability for click signals is comparable to their ability to discriminate between tonal signals of different frequencies, it suggests that dolphins can discriminate between an echo with a peak frequency of 114 kHz and an echo with a peak frequency of 118 kHz. One study found that dolphins can detect broadband signals slightly better than a pure-tone signal, but ability to discriminate between broadband sounds with different peak or center frequencies was not assessed (Au *et al.*, 2002). Assuming that dolphins' ability to discriminate temporal differences in click signals is comparable to their ability for tonal signals (Yunker and Herman, 1974), a dolphin could tell the difference between an echo that is 100 and 109 μs (using a difference limen of 0.06). Clearly, studies of dolphins' ability to discriminate between differences in the duration, bandwidth, or number of highlights in echoes are needed to make further progress on determining the salient features of echoes using the error analysis method presented here.

ACKNOWLEDGMENTS

We would like to thank Kelly Benoit-Bird, Aran Mooney, and Kim Andrews for their assistance in making the acoustic measurements of the targets. We are grateful to the trainers and staff at Epcot's Living Seas, particularly E. Putman, A. Stamper, D. Bickel, D. Clark, J. Gory, W. Fellner, M. Muraco, M. Barringer, B. Cavanaugh, L. Davis, D. Feuerbach, C. Goonen, L. Larsen-Plott, K. Odell, C. Litz, J. Davis, T. Hopkins, B. Stevens, J. Mellen, and J. Ogden. We thank Paul Nachtigall, Patricia Couvillon, Ann Peters, Nathaniel Gibbs, Kelly Benoit-Bird, and James Simmons for their constructive suggestions on earlier versions of this manuscript. Funding for this project was provided in part by grants to CMD from the American Psychological Association (Dissertation Research Award), American Association of University Women (Pacific Dissertation Fellowship), Soroptimist International (Founder Region Dissertation Fellowship), and SEASPACE (Scholarship Award) and to HEH from the Walt Disney Company's Animal Programs and New College of Florida. This publication is contribution number 1206 from the Hawaii Institute of Marine Biology.

[1]The value d' is usually calculated for a detection task, but can also be calculated for a discrimination task (match-to-sample) in which the participant is presented with a sample object and must choose a match from among three alternatives. Data collected in this three-alternative procedure for each object set were broken down into three 2×2 confusion matrices to calculate d'. In the following example, the 2×2 matrix includes only objects 1 and 2 (trials with object 3 are not counted). A hit is the choice of object 1 when object 1 was the correct match, and a false alarm is the choice of object 1 when object 2 was the correct match. d' was calculated using the proportion of hits (the number of trials in which the dolphin correctly responded that object 1 was the sample, divided by the total number of trials in which 1 was the sample) vs the proportion of false alarms (the number of trials in which the dolphin responded that object 1 was the sample, divided by the total number of trials in which object 2 was the sample). A d' value of zero indicates no discrimination behavior, and is comparable to 50% correct on a two-alternative task. A d' value of 1.0 is comparable to approximately 76% correct responses on a two-choice task.

Au, W. W. L. (**1980**). "Echolocation signals of the Atlantic bottlenose dolphin (*Tursiops truncatus*) in open waters," in *Animal Sonar Systems*, edited by R. G. Busnel and J. F. Fish (Plenum Press, New York), pp. 251–282.

Au, W. W. L. (**1993**). *The Sonar of Dolphins* (Springer, New York).

Au, W. W. L. (**2000**). "Echolocation in dolphins," in *Hearing by Whales and Dolphins*, edited by W. W. L. Au, A. N. Popper, and R. R. Fay (Springer New York), pp. 364–408.

Au, W. W. L., Andersen, L. N., Rasmussen, A. R., Roitblat, H. L., and Nachtigall, P. E. (**1995**). "Neural network modeling of a dolphin's sonar discrimination capabilities," J. Acoust. Soc. Am. **98**, 43–50.

Au, W. W. L., Floyd, R. W., Penner, R. H., and Murchison, A. E. (**1974**). "Measurement of echolocation signals of the Atlantic bottlenose dolphin, *Tursiops truncatus* Montagu, in open waters," J. Acoust. Soc. Am. **54**, 1280–1290.

Au, W. W. L., Lemonds, D. W., Vlachos, S., Nachtigall, P. E., and Roitblat, H. L. (**2002**). "Atlantic bottlenose dolphin (*Tursiops truncatus*) hearing thresholds for brief broadband signals," J. Comp. Psychol. **116**, 151–157.

Au, W. W. L., and Martin, D. (**1988**). "Sonar discrimination of metallic plates by dolphins and humans," in *Animal Sonar: Processes and Performance*, edited by P. E. Nachtigall and P. W. B. Moore (Plenum Press, New York), pp. 809–813.

Au, W. W. L., and Pawloski, J. L. (**1992**). "Cylinder wall thickness difference discrimination by an echolocating Atlantic bottlenose dolphin," J. Comp. Physiol., A **172**, 41–47.

Au, W. W. L., and Turl, C. W. (**1991**). "Material composition discrimination of cylinders at different aspect angles by an echolocating dolphin," J. Acoust. Soc. Am. **89**, 2448–2451.

Bauer, G. B., and Johnson, C. M. (**1994**). "Trained motor imitation by bottlenose dolphins (*Tursiops truncatus*)," Percept. Mot. Skills **79**, 1307–1315.

Benoit-Bird, K. J., Au, W. W. L., and Kelley, C. D. (**2003**). "Acoustic backscattering by Hawaiian lutjanid snappers. I. Target strength and swimbladder characteristics," J. Acoust. Soc. Am. **114**, 2757–2766.

Green, D. M., and Swets, J. A. (**1988**). *Signal Detection Theory and Psychophysics* (Peninsula, Los Altos, CA).

Hammer, C. E., Jr., and Au, W. W. L. (**1980**). "Porpoise echo-recognition:

An analysis of controlling target characteristics," J. Acoust. Soc. Am. **68**, 1285–1293.

Harley, H. E., Putman, E. A., and Roitblat, H. L. (**2003**). "Bottlenose dolphins perceive object features through echolocation," Nature (London) **424**, 667–669.

Herman, L. M., and Arbeit, W. R. (**1972**). "Frequency discrimination limens in the bottlenose dolphin: 1-70Ks/c," J. Auditory Res. **2**, 109–120.

McClellan, M. E., and Small, A. M. (**1965**). "Time-separation pitch associated with correlated noise burst," J. Acoust. Soc. Am. **38**, 142–143.

Nachtigall, P. E., Murchison, A. E., and Au, W. W. L. (**1980**). "Cylinder and cube discrimination by an echolocating blindfolded bottlenose dolphin," in *Animal Sonar Systems*, edited by R. G. Busnel and J. F. Fish (Plenum Press, New York), pp. 945–947.

Okimoto, G., Shizumura, R., and Lemonds, D. (**1998**). "Active biosonar systems based on multiscale signal representations and hierarchical neural networks," in *Detection and remediation technologies for mines and mine-like targets III*, Proc. SPIE **3392**, 316–323.

Roitblat, H. L., Moore, P. W. B., Nachtigall, P. E., Penner, R. H., and Au, W. W. L. (**1989**). "Natural echolocation with an artificial neural network," Int. J. Neural Networks **1**, 239–248.

Small, A. M. and McClellan, M. E. (**1963**). "Pitch associated with time delay between two pulse trains," J. Acoust. Soc. Am. **35**, 1246–1255.

Xitco, M. J., Jr. (1996). "Referential pointing by bottlenose dolphins," Southern Methodist University, Dallas (unpublished doctoral dissertation).

Xitco, M. J., Jr., and Roitblat, H. L. (**1996**). "Object recognition through eavesdropping: Passive echolocation in bottlenose dolphins," Anim. Learn Behav. **24**, 355–365.

Yunker, M. P., and Herman, L. M. (**1974**). "Discrimination of auditory temporal differences by the bottlenose dolphin and by the human," J. Acoust. Soc. Am. **56**, 1870–1875.